CLAIMS

SUB A27

1. A method of operating a filesystem, said filesystem including a live filespace accessible to users and a zombie filespace not accessible to users, said method including

recording changes to said zombie filespace in a persistent memory.

2. A method as in claim 1, including, for a deletion operation on a file in said live filespace,

transferring said file from said live filespace to said zombie filespace;

breaking links associating disk blocks with said file in a plurality of steps while said file is associated with said zombie filespace, wherein said recording of changes includes recording said breaking of links in a plurality of steps; and

altering said live filespace to reflect said deletion operation.

3. A method as in claim 1, including, for a truncation operation on a file in said live filespace,

transferring at least a portion of said file from said live filespace to said zombie filespace;

breaking links associating disk blocks with said file in a plurality of steps while a portion of said file is associated with said zombie filespace, wherein said recording of changes includes recording said breaking of links in a plurality of steps; and

altering said live filespace to reflect changes associated with said breaking of links.

4. A method as in claim 1, including, for an operation apparent to users as substantially atomic, performing said operation in a plurality of steps using said zombie filespace, wherein said recording changes is performed in said persistent memory for each of said plurality of steps.

5. A method as in claim 1, including, for an operation performed on a file having attached data elements, performing said operation using said zombie filespace.

6. A method as in claim 1, including, for an operation performed using said zombie filespace, altering a size of said zombie filespace during performance of said operation.

7. A method as in claim 1, including, for an operation performed using said zombie filespace, checkpointing said filesystem during performance of said operation.

8. A method as in claim 1, including recording changes to said live filespace in said persistent memory, wherein records of changes to said live filespace and of changes to said zombie filespace are substantially interspersed.

9. A method as in claim 1, including replaying a set of said changes in response to said record.

10. A method as in claim 1, including replaying a set of said changes to said live filespace and to said zombie filespace, wherein replay of changes includes substantial interspersed performance of changes to said live filespace and to said zombie filespace.

11. A method as in claim 1, including replaying a set of said changes in said record in response to a crash recovery by said filesystem.

12. A method as in claim 1, wherein said persistent memory includes a log of substantially all changes, within a selected time duration, to either said live filespace or said zombie filespace.

13. A method as in claim 1, wherein said persistent memory includes a log of substantially all changes, within a selected time duration, to said zombie filespace.

14. A method as in claim 1, wherein said recorded changes include a set of substantially atomic operations to said zombie filespace.

15. A method of operating a filesystem, said filesystem including a live filespace accessible to users and a zombie filespace not accessible to users, said method including

dynamically growing said zombie filespace.

16. A method as in claim 15, including, for a deletion or truncation operation on a file in said live filespace,

allocating storage within said zombie filespace for metadata associated with said file;

performing said dynamic growth in response to failure of said allocation of storage;

re-performing said allocation of storage after said dynamic growth; and

transferring said file from said live filespace to said zombie filespace.

17. A method as in claim 15, wherein said dynamic growth occurs, for an operation performed using said zombie filespace, during performance of said operation.

18. A method of operating a filesystem, said filesystem including a live filespace accessible to users and a zombie filespace not accessible to users, said method including

transfer of a file to said zombie filespace before breakage of links to blocks in said file, in response to an operation on said file, said operation using said zombie filespace.

19. A method as in claim 18, wherein, for a deletion operation on a file in said live filespace,

said transfer includes

creating a link associating said file with said zombie filespace; and

breaking a link associating said file with said live filespace;

and said deletion operation includes

breaking links associating disk blocks with said file in a plurality of steps while said file is associated with said zombie filespace, wherein said recording of changes includes recording said breaking of links in a plurality of steps; and

altering said live filespace to reflect said deletion operation.

20. A method as in claim 18, wherein, for a truncation operation on a file in said live filespace,

said transfer includes

creating a link associating at least a portion of said file with said zombie filespace; and

breaking a link associating said portion with said file in said live filespace;

and said truncation operation includes

breaking links associating disk blocks with said file in a plurality of steps while a portion of said file is associated with said zombie filespace, wherein said recording of changes includes recording said breaking of links in a plurality of steps; and

altering said live filespace to reflect changes associated with said breaking of links.

21. A method of operating a filesystem, said filesystem including a live filespace accessible to users and a zombie filespace not accessible to users, said method including

transfer of a file to said zombie filespace before performing any substantial portion of an operation on said file, said operation using said zombie filespace.

22. A method as in claim 21, wherein, for a deletion operation on a file in said live filespace,

said transfer includes

creating a link associating said file with said zombie filespace; and

breaking a link associating said file with said live filespace;

and said deletion operation includes

breaking links associating disk blocks with said file in a plurality of steps only while said file is associated with said zombie filespace, wherein said recording of changes includes recording said breaking of links in a plurality of steps; and

altering said live filespace to reflect said deletion operation.

23. A method as in claim 21, wherein, for a truncation operation on a file in said live filespace,

said transfer includes

creating a link associating at least a portion of said file with said zombie filespace; and

breaking a link associating said portion with said file in said live filespace; and said truncation operation includes

breaking links associating disk blocks with said file in a plurality of steps only while a portion of said file is associated with said zombie filespace, wherein said recording of changes includes recording said breaking of links in a plurality of steps; and

altering said live filespace to reflect changes associated with said breaking of links.

24. A method of operating a filesystem, said filesystem including a live filespace accessible to users and a zombie filespace not accessible to users, said method including

replay of an operation on a file, said operation using said zombie filespace.

25. A method as in claim 24, wherein said replay is responsive to a set of recorded changes in a persistent memory;

and including, for a deletion operation on a file in said live filespace,

transferring said file from said live filespace to said zombie filespace, and recording said transfer in said persistent memory;

breaking links associating disk blocks with said file in a plurality of steps while said file is associated with said zombie filespace, and recording said breaking of links in said persistent memory in a plurality of steps; and

altering said live filespace to reflect said deletion operation, and recording said alteration in said persistent memory.

26. A method as in claim 24, wherein said replay is responsive to a set of recorded changes in a persistent memory;

and including, for a truncation operation on a file in said live filespace,

transferring at least a portion of said file from said live filespace to said zombie filespace, and recording said transfer in said persistent memory;

breaking links associating disk blocks with said file in a plurality of steps while a portion of said file is associated with said zombie filespace, and recording said breaking of links in said persistent memory in a plurality of steps; and

altering said live filespace to reflect changes associated with said breaking of links, and recording said alteration in said persistent memory.

27. A method of operating a filesystem, said filesystem including a live filespace accessible to users and a zombie filespace not accessible to users, said method including

replay of a set of filesystem operations, said operations including at least some operations using said live filespace and at least some operations using said zombie filespace.

28. A method as in claim 27, wherein said replay is responsive to a set of recorded changes in a persistent memory;

and including, for a deletion operation on a file in said live filespace,

transferring said file from said live filespace to said zombie filespace, and recording said transfer in said persistent memory;

breaking links associating disk blocks with said file in a plurality of steps while said file is associated with said zombie filespace, and recording said breaking of links in said persistent memory in a plurality of steps; and

altering said live filespace to reflect said deletion operation, and recording said alteration in said persistent memory.

29. A method as in claim 27, wherein said replay is responsive to a set of recorded changes in a persistent memory;

and including, for a truncation operation on a file in said live filespace,

transferring at least a portion of said file from said live filespace to said zombie filespace, and recording said transfer in said persistent memory;

breaking links associating disk blocks with said file in a plurality of steps while a portion of said file is associated with said zombie filespace, and recording said breaking of links in said persistent memory in a plurality of steps; and

altering said live filespace to reflect changes associated with said breaking of links, and recording said alteration in said persistent memory.